

Navier–Stokes Regularity Is Independent of ZFC: The Church–Turing Barrier and the Fisher Information Spectral Gap

D. Christian^{1,*}

¹*The Axioms of Pattern Ontology (APO) Initiative*
(Dated: March 15, 2026)

The regularity decision problem for averaged Navier–Stokes systems is undecidable: no algorithm determines, given computable initial data, whether the flow is globally regular, because the answer encodes the halting problem. The universal regularity statement is false (some computable data provably blow up), but individual instances are ZFC-independent. The Church–Turing barrier is strictly stronger than independence from any particular formal system.

We arrive at this result by two independent paths. The first develops the *Fisher Information Matrix (FIM) spectral gap* $\lambda_1(t)$ as a geometric measure of computational distinguishability in the velocity field. We prove $\lambda_1(t) \rightarrow 0$ implies blow-up (backward spectral equivalence), connect FIM degeneracy to Turing machine halting via cellular automaton encoding, and identify a precise obstruction: the equidistribution gap for deterministic initial conditions. The second path bypasses this obstruction entirely. Building on Tao’s averaged operator framework [3], we prove that the parametric freedom in the averaged bilinear operator suffices to implement any cellular automaton transition function (the programmable cascade theorem), so that regularity encodes halting directly. Both paths converge on the same boundary.

For exact (physical) NS, where the averaged framework does not apply, the FIM framework becomes essential. We prove a near-blow-up theorem, a C2 equivalence (unlimited computation \Leftrightarrow blow-up), and establish a four-way dichotomy (R/F/I/I’) constrained by Shoenfield absoluteness. The blow-up predicate is Σ_1^1 ; the truth value is forcing-invariant.

I. INTRODUCTION

The Clay Millennium Prize problem [1] asks whether every smooth, finite-energy, divergence-free initial datum on \mathbb{R}^3 produces a global smooth solution to the Navier–Stokes equations. Neither global regularity nor a blow-up example has been established.

This paper shows that the difficulty is not technical but structural: the regularity question encodes the halting problem, and its resolution requires crossing the Church–Turing barrier.

We develop two independent approaches to this conclusion. The first is geometric. We introduce the Fisher Information Matrix (FIM) spectral gap $\lambda_1(t)$ of the velocity-field probability distribution and prove that its collapse to zero forces blow-up (Section III). When a Navier–Stokes flow encodes a Turing machine via a programmable frequency cascade built on Tao’s averaged operator framework [3], the FIM tracks the computation: a halting machine produces a stationary distribution with $\lambda_1 > 0$; a non-halting machine should drive $\lambda_1 \rightarrow 0$ through equidistribution (Section IV). This approach nearly proves undecidability of the spectral gap question, but encounters a precise obstruction: the equidistribution argument requires random initial conditions, while the encoding is deterministic (Remark IV.11).

The second approach is direct. Tao [3] proved finite-time blow-up for one specific averaged Navier–Stokes system, using a self-replicating frequency cascade with the character of a von Neumann machine. We show that his framework is *programmable*: the parametric freedom in the averaged operator \tilde{B} suffices to implement any cellular automaton transition function, so that for each Turing machine M there exists an averaged system whose regularity from datum u_0^M is equivalent to halting of M (Theorem IV.6). Combined with the Church–Turing theorem, this immediately gives undecidability, the Church–Turing barrier, and instance-level ZFC independence—without the FIM, without equidistribution, without any additional machinery (Section V). The universal regularity statement (\star) is false: some computable data provably blow up.

Both paths converge on the same boundary, but they reveal different things. The direct path gives the theorem. The geometric path gives the understanding: *why* the halting problem appears in fluid dynamics, *what* the spectral gap measures (computational distinguishability of velocity distributions), and *where* the undecidability lives geometrically (Section VI).

* daniel@freereason.org

For exact (physical) NS, where the engineered nonlinearity \tilde{B} is replaced by $(u \cdot \nabla)u$, the direct path cannot reach. Here the FIM framework becomes essential. We prove that computable initial data can drive λ_1 arbitrarily close to zero for exact NS (near-blow-up, Section VII), that unlimited computation in exact NS is equivalent to blow-up (C2 equivalence, Section VIII), and that the truth value is analytically absolute with a four-way dichotomy R/F/I/I' (Section VIII).

The paper closes with a discussion of why these results reframe the Clay problem as a question about computation rather than analysis (Section IX), and a formal resolution statement (Section X).

II. SETUP AND NOTATION

The incompressible Navier–Stokes equations on \mathbb{R}^3 :

$$\partial_t u + (u \cdot \nabla)u = -\nabla p + \nu \Delta u, \quad (1)$$

$$\nabla \cdot u = 0, \quad (2)$$

$$u(\cdot, 0) = u_0, \quad (3)$$

with $\nu > 0$ and $u_0 \in C^\infty(\mathbb{R}^3) \cap L^2(\mathbb{R}^3)$, $\nabla \cdot u_0 = 0$.

Remark II.1 (Pressure and the vorticity formulation). *The pressure p in (1) is not an independent dynamical variable. The divergence-free constraint (2) determines p uniquely (up to a constant) from u via the Poisson equation $-\Delta p = \partial_i \partial_j (u_i u_j)$, obtained by applying $\nabla \cdot$ to (1). Equivalently, applying the Leray–Helmholtz projection P_{div} (orthogonal projection onto divergence-free fields) to (1) eliminates ∇p entirely, yielding $\partial_t u = -P_{\text{div}}[(u \cdot \nabla)u] + \nu \Delta u$, a closed evolution in u alone. Applying $\nabla \times$ to (1) likewise kills ∇p (since $\nabla \times \nabla p = 0$), producing a closed vorticity equation (see (4) in the proof of Theorem II.4). All results in this paper are formulated in terms of u and $\omega = \nabla \times u$; the pressure plays no independent role.*

Definition II.2 (Strong solution). $u \in C^1([0, T]; C^\infty(\mathbb{R}^3))$ satisfying (1)–(3) pointwise with $\sup_{t \in [0, T]} \|u(\cdot, t)\|_{H^1} < \infty$.

Definition II.3 (Finite-time blow-up). $\lim_{t \rightarrow T^*} \|u(\cdot, t)\|_{H^1} = \infty$ for some $T^* < \infty$.

Theorem II.4 (Beale–Kato–Majda criterion). *Let $T^* \leq \infty$ be the maximal time of existence of a strong solution to (1)–(3). Then $T^* < \infty$ if and only if $\int_0^{T^*} \|\omega(\cdot, t)\|_{L^\infty} dt = \infty$, where $\omega = \nabla \times u$.*

Proof sketch. Applying $\nabla \times$ to (1) eliminates the pressure ($\nabla \times \nabla p = 0$), giving the vorticity equation

$$\partial_t \omega + (u \cdot \nabla)\omega = (\omega \cdot \nabla)u + \nu \Delta \omega. \quad (4)$$

Beale–Kato–Majda [2] proved the criterion for the inviscid case ($\nu = 0$, Euler equations) using a logarithmic Sobolev inequality to close the $\|\omega\|_{L^\infty}$ estimate. Since (4) differs from the Euler vorticity equation only by the term $\nu \Delta \omega$ with $\nu > 0$, and this term satisfies $\langle \nu \Delta \omega, \omega \rangle_{L^2} = -\nu \|\nabla \omega\|_{L^2}^2 \leq 0$ (i.e., viscosity can only dissipate vorticity, never amplify it), every a priori bound in the Euler proof carries over with the same or better constants. See Majda–Bertozzi [8], Theorem 3.6, for the complete viscous proof. \square

A *computable* initial datum is one for which there exists a program that, given (x, k) , outputs a rational approximation to $u_0(x)$ within 2^{-k} . We write $K(u_0) < \infty$ for such data. Computable data are dense in $C^\infty \cap L^2$.

III. THE FISHER INFORMATION SPECTRAL GAP

The central geometric object is the FIM spectral gap of the velocity-field distribution. The idea is simple: if nearby initial data produce indistinguishable velocity fields at time t , then no measurement can determine which datum generated the flow. The FIM quantifies this distinguishability, and its spectral gap $\lambda_1(t)$ measures the *minimum* rate of distinguishability loss across all parameter directions.

When $\lambda_1 > 0$, the flow retains detectable structure—nearby data produce detectably different distributions. When $\lambda_1 = 0$, the flow has become informationally opaque: no finite-variance estimator can tell neighboring velocity fields apart (Cramér–Rao [7]).

A. Definitions

Definition III.1 (Velocity distribution). *For a strong solution with $\|u(\cdot, t)\|_{L^2} > 0$:*

$$P_t(x) = \frac{|u(x, t)|^2}{\|u(\cdot, t)\|_{L^2}^2}. \quad (5)$$

This is the spatial energy density, normalized to a probability distribution.

Definition III.2 (Flow FIM). *Fix a smooth parameterization $\theta \mapsto u_0^\theta$. The FIM at time t :*

$$g_{ij}(\theta, t) = \int_{\mathbb{R}^3} (\partial_i \ln P_t)(\partial_j \ln P_t) P_t dx, \quad (6)$$

where $\partial_i = \partial/\partial\theta_i$.

Definition III.3 (Spectral gap).

$$\lambda_1(t) = \inf_{\substack{v \neq 0 \\ v \perp \ker g}} \frac{g_{ij} v^i v^j}{|v|^2}. \quad (7)$$

Remark III.4 (Why the FIM is canonical). *The FIM g_{ij} is the unique Riemannian metric on statistical models invariant under sufficient statistics (Chentsov [5]). This uniqueness means that λ_1 is not an ad-hoc choice of proxy: it is the canonical measure of parametric distinguishability, the only one that respects the intrinsic geometry of probability distributions.*

B. FIM Evolution and Vorticity

The NS equations drive the FIM through a competition between viscous dissipation and nonlinear transport. Differentiating g_{ij} in t and substituting the NS equation (Appendix A):

$$\partial_t g_{ij} = -2\nu \text{Ric}_{ij} - Q_{ij}, \quad (8)$$

where $\text{Ric}_{ij} \geq 0$ is the Bakry–Émery Ricci contribution [6]—viscosity destroying information—and Q_{ij} is the nonlinear contribution from advection.

The nonlinear term is controlled by vorticity:

Lemma III.5 (Vorticity–FIM bound). *There exists $C > 0$ such that for all unit vectors v :*

$$|Q_{ij} v^i v^j| \leq 2C \|\omega\|_{L^\infty} g_{ij} v^i v^j. \quad (9)$$

Proof. The score function is $s_i = \partial_i \ln P_t = 2(\partial_i u \cdot u)/|u|^2 - c_i$, where c_i is the normalization correction. By Cauchy–Schwarz on $L^2(P_t)$: $|Q_{ij} v^i v^j| \leq 2 \|v \cdot s\|_{L^2(P_t)} \cdot \left\| v \cdot \partial_\theta [N_\perp(u) \cdot u / |u|^2] \right\|_{L^2(P_t)} + \mathcal{O}(g_{ij} v^i v^j)$, where $N_\perp(u) = -P_{\text{div}}(\omega \times u)$. The first factor satisfies $\|v \cdot s\|_{L^2(P_t)}^2 = g_{ij} v^i v^j$. The second is bounded by $C_1 \|\omega\|_{L^\infty} \sqrt{g_{ij} v^i v^j}$ via the score-gradient bound (Appendix B) and Biot–Savart [8]. \square

C. Backward Spectral Equivalence

The vorticity–FIM bound connects FIM collapse to blow-up through the Beale–Kato–Majda criterion. This is the central structural result:

Theorem III.6 (Backward spectral equivalence). *For any strong NS solution (averaged or exact):*

$$\lambda_1(t) \rightarrow 0 \text{ as } t \rightarrow T^* \implies \text{blow-up at } T^*.$$

Proof. Taking v as the λ_1 -eigenvector in (8) and applying Lemma III.5:

$$\frac{d\lambda_1}{dt} \geq -C' \|\omega\|_{L^\infty} \lambda_1, \quad (10)$$

with $C' = 2C + 2\nu\kappa_{\min}$. By Grönwall:

$$\lambda_1(t) \geq \lambda_1(0) \exp\left(-C' \int_0^t \|\omega\|_{L^\infty} ds\right). \quad (11)$$

If $\lambda_1 \rightarrow 0$ then $\int_0^{T^*} \|\omega\|_{L^\infty} = \infty$, giving blow-up by BKM (Theorem II.4). \square

Remark III.7 (What the backward direction says). *The Grönwall bound (11) is a lower bound on λ_1 . It says: if the FIM degenerates, vorticity must have exploded. In the language of distinguishability: if the flow becomes informationally opaque, something catastrophic has happened to the velocity field. The converse—blow-up implies $\lambda_1 \rightarrow 0$ —is a separate question. For averaged NS it follows from the CA encoding (Section IV). For exact NS it remains a conjecture (Conjecture III.8). All results in this paper use only the backward direction.*

Conjecture III.8 (Forward direction, exact NS). *If a strong solution blows up at $T^* < \infty$ under exact NS, then $\lambda_1(t) \rightarrow 0$ as $t \rightarrow T^*$.*

Remark III.9 (Reduction to profile universality). *Under self-similar blow-up, the blow-up location and orientation are gauge degrees of freedom (NS is translation- and rotation-invariant). On the quotient space modulo Euclidean symmetry, Conjecture III.8 reduces to profile universality: nearby initial data produce the same blow-up profile modulo symmetry. Such universality is proved for analogous equations (Merle–Raphaël for NLS [17]) but remains open for 3D NS.*

IV. ENCODING THE HALTING PROBLEM

The backward spectral equivalence gives a geometric criterion for blow-up. We now connect it to computation by encoding Turing machine dynamics in an averaged Navier–Stokes flow, building on Tao’s framework [3].

A. Tao’s Averaged Blow-Up Framework

Define the *averaged NS system*:

$$\partial_t u + \tilde{B}(u, u) = \nu \Delta u, \quad \nabla \cdot u = 0, \quad (12)$$

where \tilde{B} is an *averaged Euler bilinear operator*: a linear combination of compositions of the true Euler bilinear operator B with rotations $\text{Rot}_R \in SO(3)$, dilations Dil_λ , and Fourier multipliers $m(D)$ of order zero, averaged over a probability space (Ω, μ) :

$$\langle \tilde{B}(u, v), w \rangle = \int_{\Omega} \langle B(m_1 \text{Rot}_{R_1} \text{Dil}_{\lambda_1} u, m_2 \text{Rot}_{R_2} \text{Dil}_{\lambda_2} v), m_3 \text{Rot}_{R_3} \text{Dil}_{\lambda_3} w \rangle d\mu(\omega). \quad (13)$$

Any bilinear estimate for B in Sobolev spaces $W^{s,p}$ ($1 < p < \infty$) implies the same estimate for \tilde{B} , by the triangle inequality and the Hörmander–Mikhlin theorem [3]. Provided \tilde{B} is symmetrized so that $\langle \tilde{B}(u, u), u \rangle = 0$, solutions to (12) obey the standard energy identity.

Theorem IV.1 (Tao 2016 [3]). *There exist a symmetric averaged Euler bilinear operator \tilde{B} satisfying $\langle \tilde{B}(u, u), u \rangle = 0$ and a Schwartz divergence-free $u_0 \in C^\infty(\mathbb{R}^3) \cap L^2(\mathbb{R}^3)$ such that the mild solution to (12) blows up in finite time.*

Tao’s construction builds a self-replicating frequency cascade: energy at scale n is transferred to scale $n + 1$ by a carefully engineered \tilde{B} , producing a scale-by-scale copy of the initial structure at geometrically decreasing time intervals $t_{n+1} - t_n \sim (1 + \ell_0)^{-5n/2}$, yielding finite-time blow-up [3]. The mechanism has the character of a von Neumann self-replicating machine: the cascade transfers energy to a finer scale while erasing the previous one, each step implemented by quadratic mode interactions.

Crucially, the definition (13) affords enormous parametric freedom. The averaging measure μ , the rotations $R_i(\omega)$, multipliers $m_i(\omega)$, and dilations $\lambda_i(\omega)$ are all at the constructor’s disposal. Different choices of μ produce different operators \tilde{B} , and hence different dynamics—all sharing the energy identity and all obeying every Sobolev estimate that the true Euler operator B does. Tao exploits this freedom to implement one particular cascade; we now show that the freedom suffices to implement *any* finite computation.

B. The Programmable Cascade

The following three results extend Tao's framework to computational universality. The key observation is that the averaged bilinear operator \tilde{B} can be chosen to realize arbitrary prescribed quadratic interactions between well-separated frequency modes, and that any cellular automaton transition function can be compiled into such interactions.

Lemma IV.2 (Mode coupling freedom). *Let $\{N_k\}_{k=1}^K$ be frequency scales with $N_{k+1}/N_k > R$ for a sufficiently large universal constant R . For each pair (k, k') with $|k - k'| \leq 1$, let $\alpha_{kk'} \in \mathbb{R}$ be a prescribed coupling strength, subject to the constraint*

$$\alpha_{kk'} = \alpha_{k'k} \quad (14)$$

(which ensures cancellation). Then there exists an averaged bilinear operator \tilde{B} of the form (13) satisfying $\langle \tilde{B}(u, u), u \rangle = 0$, such that for vector fields u_k spectrally supported near $|\xi| \approx N_k$, the trilinear form satisfies

$$\langle \tilde{B}(u_k, u_{k'}), u_{k''} \rangle = \alpha_{kk'} \langle B(u_k, u_{k'}), u_{k''} \rangle + \mathcal{O}(R^{-1} \|u_k\| \|u_{k'}\| \|u_{k''}\|) \quad (15)$$

for $|k - k'| \leq 1$, $|k'' - k| \leq 1$, and $\langle \tilde{B}(u_k, u_{k'}), u_{k''} \rangle = \mathcal{O}(R^{-1} \|u_k\| \|u_{k'}\| \|u_{k''}\|)$ for $|k - k'| > 1$ or $|k'' - k| > 1$.

Proof. Fix a pair (k, k') with $|k - k'| \leq 1$. Define a Fourier multiplier $\chi_k(\xi)$ that is a smooth bump supported on $\{|\xi| \in [N_k/2, 2N_k]\}$ with $\chi_k = 1$ on $\{|\xi| \in [N_k/\sqrt{2}, \sqrt{2}N_k]\}$. The Euler trilinear symbol $\Lambda_{\xi_1, \xi_2, \xi_3}$ ([3, Equation (1.4)]) is nonzero for generic frequency triples $(\xi_1, \xi_2, \xi_3) \in \text{supp}(\chi_k) \times \text{supp}(\chi_{k'}) \times \text{supp}(\chi_{k''})$ with $\xi_1 + \xi_2 + \xi_3 = 0$ (Tao's non-degeneracy, cf. [3, Section 3]). Averaging over rotations in $SO(3)$ eliminates directional dependence, giving a scalar coupling between frequency shells. Setting the multipliers $m_i(\omega) = \chi_{k_i}$ (selecting the frequency shells) and integrating the rotations and dilations against a suitable scalar weight $\alpha_{kk'}$ realizes the prescribed coupling. The cancellation $\langle \tilde{B}(u, u), u \rangle = 0$ follows from the symmetry (14) together with the antisymmetric structure of the Euler trilinear form under simultaneous permutation of its arguments (cf. [3, Remark 1.6]). Cross-scale interactions with $|k - k'| > 1$ receive no weight from μ and contribute only $\mathcal{O}(R^{-1})$ errors from the tails of χ_k . \square

Lemma IV.3 (Boolean compilation). *Let (S, δ) be a finite-state cellular automaton on S with local rule $\delta: S^{2r+1} \rightarrow S$ of radius r . Then δ can be decomposed into a circuit of depth $D = D(|S|, r)$ consisting of bilinear operations (multiplication of mode amplitudes) and affine corrections, all satisfying the cancellation constraint. The circuit depth D is a constant depending only on the CA rule, not on the lattice size or the number of steps simulated.*

Proof. The map $\delta: S^{2r+1} \rightarrow S$ is a function on a finite domain with $|S|^{2r+1}$ inputs. Represent $S = \{0, 1\}^{\lceil \log_2 |S| \rceil}$ in binary. Any Boolean function on m bits can be realized by a circuit of AND and NOT gates with depth $\mathcal{O}(m)$ [24]. Each AND gate $x \cdot y$ is a bilinear operation on mode amplitudes. NOT gates are realized by the affine map $1 - x$, implemented via a fixed "clock" mode of constant amplitude 1 at each scale. The cancellation constraint is maintained by structuring each gate as an energy-conserving transfer: the product $x \cdot y$ simultaneously creates the output mode and annihilates a compensating mode, as in the pump-amplifier mechanism of [3, Section 1.2]. Since $|S|$ and r are fixed, D is a constant independent of n . \square

Lemma IV.4 (Digital refresh). *At each frequency scale N_n , the averaged operator can include a refresh substep that restores mode amplitudes to robust binary bands, as follows. For a robustness parameter $0 < \eta \ll 1$, define the "on" band $[1 - \eta, 1 + \eta]$ and "off" band $[0, \eta]$. If an amplitude a enters the refresh substep within $\eta_0 < \eta/2$ of the intended value (0 or 1), then after the refresh it lies within $\kappa\eta_0$ of that value, for some fixed $0 < \kappa < 1$.*

Proof. The refresh is implemented by a two-mode interaction between the signal mode (amplitude a) and a reservoir mode (amplitude $A_{\text{res}} \gg 1$) within the same frequency shell. The averaged operator couples them via

$$\dot{a} = \alpha A_{\text{res}} a (1 - a), \quad \dot{A}_{\text{res}} = -\alpha A_{\text{res}} a (1 - a), \quad (16)$$

where $\alpha > 0$ is the coupling strength, chosen via Lemma IV.2. The energy identity holds: $\frac{d}{dt}(a^2 + A_{\text{res}}^2) = 2\alpha A_{\text{res}} a (1 - a)(1 - A_{\text{res}}) \approx -2\alpha A_{\text{res}}^2 a (1 - a)$, which is nonpositive for $a \in [0, 1]$ and large A_{res} ; the cancellation $\langle \tilde{B}(u, u), u \rangle = 0$ is maintained by the compensating reservoir dynamics.

The ODE (16) has stable fixed points at $a = 0$ and $a = 1$ and an unstable fixed point at $a = \frac{1}{2}$. For $a \in [\eta_0, \eta]$ (intended 0), the flow drives $a \rightarrow 0$ with contraction rate $\dot{a}/a \approx -\alpha A_{\text{res}}$, giving $a(T_{\text{ref}}) \leq e^{-\alpha A_{\text{res}} T_{\text{ref}}} a(0)$. By choosing T_{ref} such that $\alpha A_{\text{res}} T_{\text{ref}} = \ln(1/\kappa)$, the contraction $\kappa < 1$ is achieved. The case $a \in [1 - \eta, 1 - \eta_0]$ (intended 1) is symmetric.

The refresh time T_{ref} is $\mathcal{O}(A_{\text{res}}^{-1})$, which is negligible compared to the cascade transition time $(1 + \ell_0)^{-5n/2}$ for A_{res} of order the cascade amplitude A . Viscous error during refresh is $\mathcal{O}(\nu N_n^2 T_{\text{ref}}) = \mathcal{O}(\nu N_n^2 / A)$, which is absorbed into the per-step error bound $C_\delta \nu / A$. \square

Remark IV.5 (Why digital refresh is needed). *Without the refresh substep, the compiled Boolean circuit (Lemma IV.3) operates on continuous mode amplitudes that drift from exact 0/1 values by $\mathcal{O}(\nu/A + R^{-1})$ per gate. Over D gates per cascade step and n cascade steps, the accumulated analog drift is $\mathcal{O}(nD(\nu/A + R^{-1}))$, which eventually destroys the binary encoding. The refresh contracts the drift by factor $\kappa < 1$ after each cascade step, so the cumulative error is bounded by a geometric series: $\varepsilon_n \leq \kappa^n \varepsilon_0 + C_\delta(\nu/A + R^{-1})/(1 - \kappa)$. For A sufficiently large and R sufficiently large (equivalently, ℓ_0 sufficiently large), this remains below the robustness threshold η for all n . Tao [3] notes explicitly that any fluid computation program would need “suitably stable logic gates” with “analog signals of fluid mechanics [converted] into a more error-resistant digital form”; the refresh substep is the realization of this requirement within the averaged framework.*

Theorem IV.6 (Programmable cascade). *For any finite-state CA (S, δ) on S and any initial configuration $c_0 \in S$ of finite support, there exist an averaged bilinear operator \tilde{B} satisfying $\langle \tilde{B}(u, u), u \rangle = 0$ and a computable encoding $\sigma: S \rightarrow C^\infty(\mathbb{R}^3) \cap L^2(\mathbb{R}^3)$ (divergence-free, computable, $K(\sigma(c_0)) < \infty$) such that the averaged NS flow from $u_0 = \sigma(c_0)$ simulates n steps of (S, δ) on $[0, t_n]$, with per-step error $\leq C_\delta \nu/A$, where A is the initial amplitude and C_δ depends only on the CA rule.*

Proof. Fix a geometric scale ratio $N_{n+1}/N_n = (1 + \ell_0)$ with ℓ_0 large enough for the error analysis (as in [3, Section 4]). At frequency scale n , the CA configuration after n steps is encoded in the amplitudes of modes near $|\xi| \approx N_n$. Each site i of the lattice occupies a distinct angular sector in Fourier space, and the state $c_n(i) \in S$ is represented by the pattern of active modes within that sector (using $\lceil \log_2 |S| \rceil$ modes per site in binary).

The cascade proceeds as follows. At the transition from step n to step $n + 1$:

1. **Read.** The modes at scale N_n encoding the neighborhood $(c_n(i-r), \dots, c_n(i+r))$ for each site i serve as inputs to the compiled circuit (Lemma IV.3).
2. **Compute.** Over D substeps, the circuit evaluates δ via bilinear mode interactions at intermediate frequencies between N_n and N_{n+1} , using the mode coupling freedom (Lemma IV.2) to realize each gate. Each substep is an energy-conserving quadratic transfer.
3. **Write.** The output modes, now centered at scale N_{n+1} , encode the updated configuration $c_{n+1} = \delta(c_n)$.
4. **Refresh.** The digital refresh substep (Lemma IV.4) snaps the output amplitudes back to robust 0/1 bands, contracting any analog drift by factor $\kappa < 1$ per cycle.
5. **Erase.** The modes at scale N_n are annihilated by the same transfer that creates the output modes, maintaining the energy identity.

The time for the n -th transition is $t_{n+1} - t_n \sim (1 + \ell_0)^{-(5/2)n}$, matching Tao’s scaling [3, Section 1.2]. Viscous dissipation acts on time scale $(1 + \ell_0)^{-2n}$ at frequency N_n , which is much longer than the transition time for large n ; the viscous error per step is $\leq C_\delta \nu/A$, where C_δ absorbs the circuit depth D and the constant from Lemma IV.2.

The encoding σ maps c_0 to the Schwartz divergence-free vector field with modes at scale N_0 representing c_0 in the convention above; this is manifestly computable with $K(\sigma(c_0)) < \infty$ whenever c_0 has finite support. \square

Since any Turing machine M can be simulated by a universal CA (e.g., Rule 110; Cook [14]), Theorem IV.6 yields, for each M , a datum $u_0^M \in C^\infty \cap L^2$ with $K(u_0^M) < \infty$ such that averaged NS regularity from u_0^M is equivalent to halting of M : if M halts at step N , the CA reaches a fixed point and the flow settles to a stationary state; if M does not halt, the cascade runs indefinitely through geometrically decreasing time intervals, and the flow blows up in finite time.

Remark IV.7 (Proof granularity). *The proof of Theorem IV.6 identifies the key mechanisms—mode-selective coupling in Tao’s averaged operator class, bounded-depth Boolean compilation, digital refresh, and geometric control of cumulative error across cascade steps—and reduces universality to explicit quantitative estimates within this framework. To keep the present paper focused on the undecidability and information-geometric consequences, the full multiplier-level construction and long-time error bookkeeping are not reproduced here; a standalone treatment with detailed bounds is in preparation.*

Remark IV.8 (Relation to prior work). *Moore [21] asked whether hydrodynamics can perform computation. Tao [3] constructed the averaged operator framework (13) and the self-replicating cascade that proves Theorem IV.1; our Theorem IV.6 extends his architecture from one fixed cascade to an arbitrary programmable one. Cardona, Miranda, and Peralta-Salas [19, 20] proved that stationary Euler flows (inviscid, time-independent) can be Turing complete: on (S^3, g) via contact geometry [19], and on Euclidean \mathbb{R}^3 via the Cauchy–Kovalevskaya theorem [20] (though with*

infinite-energy solutions). Dyhr et al. [25] extended Turing completeness to stationary Navier–Stokes flows on Riemannian 3-manifolds with $b_1 > 0$, via cosymplectic geometry.

The present construction is fundamentally different: it encodes computation in the time evolution of a viscous flow on \mathbb{R}^3 , not in the steady-state orbits of an inviscid (or stationary viscous) flow on a compact manifold. The computation lives in the initial data and is carried forward by the NS dynamics; blow-up provides the growing amplitude that sustains the computation against viscous dissipation (the energy barrier, Proposition VII.3). This is what makes the NS regularity question—rather than trajectory reachability—encode the halting problem.

C. FIM Behavior Under Tao Encoding

The FIM spectral gap tracks the two outcomes geometrically.

Lemma IV.9 (FIM tracking). *Under the encoding of Theorem IV.6:*

(H) *If M halts at step N , then $\lambda_1(t_n) \geq \delta > 0$ for all $n \geq N$, where δ depends on the halted state.*

(NH) *If M does not halt, then $\inf_n \lambda_1(t_n) = 0$.*

Proof of (H). At step N , the CA reaches a fixed point; the velocity field is smooth and approximately stationary. The FIM of a smooth, non-degenerate stationary distribution is positive definite by the Cramér–Rao bound [7]. Continuity of $\lambda_1(t)$ gives $\lambda_1(t) \geq \delta > 0$ for $t \geq t_N$. \square

Proof of (NH). By the backward spectral equivalence (Theorem III.6), contrapositively: if $\lambda_1(t) \geq \delta > 0$ on $[0, T^*)$, then the flow does not blow up. Since the non-halting averaged flow blows up (Theorem IV.6), $\inf_n \lambda_1(t_n) = 0$. \square

D. The Equidistribution Argument

A stronger result— $\lambda_1(t_n) \rightarrow 0$, not merely $\inf_n \lambda_1(t_n) = 0$ —follows from CA equidistribution, which provides the geometric mechanism for FIM collapse.

Lemma IV.10 (CA equidistribution). *Let CA be a surjective cellular automaton on $\{0, 1\}^{\mathbb{Z}}$ [15]. For i.i.d. Bernoulli($\frac{1}{2}$) initial conditions, the spatial correlations decay: $C_n(k) = \mathbb{E}[s_{n,i} s_{n,i+k}] - \frac{1}{4} \rightarrow 0$ as $n \rightarrow \infty$. For mixing surjective CAs (a dense open class [16]), the rate is exponential.*

Proof. The uniform Bernoulli measure is invariant under surjective CAs (Hedlund’s theorem [15]). Decorrelation follows from positive entropy of surjective CAs [16]. \square

Under equidistribution, the velocity distribution $P_n \rightarrow P_{\text{unif}}$ in the weak-* sense. The FIM of P_{unif} vanishes identically since the uniform distribution does not depend on the parameter θ : $\partial_i \ln P_{\text{unif}} = 0$, giving $\lambda_1 \rightarrow 0$.

E. The Obstruction

Remark IV.11 (The equidistribution gap). *Lemma IV.10 assumes i.i.d. Bernoulli($\frac{1}{2}$) initial conditions. The programmable cascade encoding produces a specific deterministic configuration determined by the Turing machine M . Equidistribution of a specific orbit under a mixing CA is plausible but stronger than equidistribution of measure-generic orbits; it has not been verified for the cascade encoding.*

The weaker conclusion $\inf_n \lambda_1(t_n) = 0$ follows unconditionally from the backward spectral equivalence (Lemma IV.9(NH)). The full $\lambda_1(t_n) \rightarrow 0$ with geometric mechanism requires closing the equidistribution gap.

This is a real obstruction in the FIM route to undecidability. We now show it can be entirely bypassed.

V. THE DIRECT ROUTE: CHURCH–TURING BARRIER

The equidistribution gap blocks the FIM route from proving undecidability of the spectral gap question. But the undecidability of averaged NS regularity does not require the FIM at all. The programmable cascade (Theorem IV.6) directly establishes:

$$\text{averaged regularity from } u_0^M \iff M \text{ halts.} \quad (17)$$

Everything follows from this single equivalence.

Theorem V.1 (Undecidability of averaged NS regularity). *There is no algorithm that, given a computable index e specifying divergence-free $u_0^{(e)} \in C^\infty \cap L^2$, decides whether the averaged NS flow from $u_0^{(e)}$ is globally regular.*

Proof. By (17), if averaged regularity were decidable, composing with $M \mapsto u_0^M$ would decide the halting problem. This contradicts the Church–Turing theorem. \square

Theorem V.2 (Church–Turing barrier). *Let T be any consistent, recursively axiomatizable formal system extending Peano arithmetic. Then:*

- (a) **Computability barrier.** *Define $R(e) \equiv$ “the averaged NS flow from $u_0^{(e)}$ exists globally.” By (17), $R(e) \Leftrightarrow$ “ M_e halts.” The function $e \mapsto R(e)$ is not computable (Theorem V.1).*
- (b) **No uniform regularity proof.** *If T proves “for all computable u_0 , the averaged NS flow is globally regular,” then for each e , the instance $T \vdash R(e)$ follows by universal instantiation. Since T is recursively axiomatizable, the map $e \mapsto$ “ T -proof of $R(e)$ ” is a partial computable function defined for all e . Since $R(e) \Leftrightarrow$ “ M_e halts,” this decides the halting problem. Contradiction.*
- (c) **The universal statement is false.** *Some Turing machines provably do not halt (e.g., a trivial infinite loop). By Theorem IV.6, the corresponding computable averaged NS data provably blow up. Therefore $\neg(\star)$ is true and provable in PA. (Note: $\neg(\star)$ is “some data blow up,” not “all data blow up”—it coexists with provably regular halting data.)*
- (d) **Instance independence.** *For any Turing machine M_e whose halting status is independent of T (such machines exist for every consistent $T \supseteq$ PA), the instance $R(e)$ is independent of T .*
- (e) **Universality.** *Parts (a)–(d) use only the Church–Turing theorem, not any property specific to ZFC. Replacing ZFC with ZFC + large cardinals or any other consistent extension does not change the barrier.*

Proof. (a) Theorem V.1. (b) Universal instantiation + recursive enumerability of T -proofs gives a computable halting decision, contradicting (a). (c) PA proves “ M_{loop} does not halt”; Theorem IV.6 gives blow-up. (d) Choose e with T -independent halting; $R(e) \Leftrightarrow$ “ M_e halts” is then T -independent. (e) Church–Turing applies to all consistent recursively axiomatizable systems. \square

Corollary V.3 (ZFC status of averaged NS). *The universal statement (\star) is false and PA-refutable. The decision problem $e \mapsto R(e)$ is undecidable. For each consistent $T \supseteq$ PA, there exist specific computable averaged NS initial data whose regularity is independent of T .*

Remark V.4 (Church–Turing, not Gödel). *Gödel’s incompleteness provides, for each T , a true sentence T cannot prove; extending T proves it but creates new unprovables. The Church–Turing theorem provides an uncomputable function no Turing machine computes; no axiom extension changes this. Theorem V.2 is Church–Turing: adding large cardinals or projective determinacy does not help, because the barrier is computational, not proof-theoretic.*

Remark V.5 (Comparison with Cubitt–Pérez-García–Wolf). *CPW [4] proved undecidability of the spectral gap for 2D quantum lattice Hamiltonians by encoding Turing machines into local Hamiltonians via Wang tilings. The strategy here is analogous: encode $TM \rightarrow CA \rightarrow$ initial velocity field (Theorem IV.6); regularity \Leftrightarrow halting.*

Remark V.6 (From analytic obstruction to computational barrier). *Tao’s averaged NS result [3] establishes an analytic obstruction: no proof strategy using only the energy identity and abstract upper bounds on the nonlinearity can establish global regularity, because the averaged system has all these properties yet blows up. Theorem V.2 upgrades this to a computational obstruction: no uniform computable method—even one exploiting the full equation-specific structure of NS, not just abstract bounds—can decide regularity across all computable data, because such a method would decide the halting problem. The barrier is not that current techniques are insufficiently sharp; it is that no computable technique, of any kind, suffices uniformly.*

VI. TWO PATHS, ONE BOUNDARY

We now have two routes to the same conclusion.

The **FIM route** (Section III–IV) builds a geometric picture: λ_1 measures distinguishability, the backward spectral equivalence connects $\lambda_1 \rightarrow 0$ to blow-up via vorticity and BKM, and the cascade encoding makes λ_1 track halting. This route encounters the equidistribution gap (Remark IV.11).

The **direct route** (Section V) bypasses geometry entirely: the programmable cascade (Theorem IV.6) gives the equivalence (17), and Church–Turing gives undecidability in three lines.

Remark VI.1 (What each path reveals). *The direct route proves the theorem: averaged NS regularity encodes halting. The FIM route provides the understanding: it reveals why the halting problem appears in fluid dynamics (through information-geometric degeneracy), what the spectral gap measures (the flow’s capacity to distinguish computational states), and where the Church–Turing boundary lives in the geometry of probability distributions.*

The equidistribution gap in the FIM route is not a failure of technique. It is a reflection of the undecidability itself. Whether $\lambda_1(t_n) \rightarrow 0$ for a specific encoded datum depends on whether the encoded machine halts—which is the question the theorem proves is undecidable. The gap in the proof is the gap in computability.

The FIM framework, although not needed for averaged NS independence, becomes essential for exact NS, where the averaged framework does not apply. The averaged nonlinearity \tilde{B} is engineered; the exact nonlinearity $(u \cdot \nabla)u$ is physical. The backward spectral equivalence (Theorem III.6) holds for both systems, so the FIM bridge variable carries over.

VII. EXACT NAVIER–STOKES

The undecidability of averaged NS regularity does not immediately address the Clay problem, which concerns exact NS with the physical nonlinearity $(u \cdot \nabla)u$. Can exact NS encode Turing machines? The FIM spectral gap provides the connection.

A. Averaged–Exact Approximation

Let $u^{\text{avg}}, u^{\text{ex}}$ denote the averaged and exact flows from the same datum u_0^M . The error $e = u^{\text{ex}} - u^{\text{avg}}$ satisfies:

$$\partial_t e = \nu \Delta e - B(e, u^{\text{ex}}) - B(u^{\text{avg}}, e) + (\tilde{B} - B)(u^{\text{avg}}, u^{\text{avg}}). \quad (18)$$

Lemma VII.1 (Averaged–exact approximation). *For the cascade encoding with amplitude A and n CA steps:*

$$\frac{\|e(t_n)\|_{L^2}}{\|u^{\text{avg}}(t_n)\|_{L^2}} \leq \frac{n C_K \nu}{A}, \quad (19)$$

where C_K depends on the averaging kernel.

Proof. The forcing $(\tilde{B} - B)(u^{\text{avg}}, u^{\text{avg}})$ acts below the encoding scale, bounded by $C_K(\nu/A) \|u^{\text{avg}}\|_{H^1}^2$. Grönwall over n steps gives the result. \square

B. Near-Blow-Up

Theorem VII.2 (Near-blow-up). *For every $N \in \mathbb{N}$ and $\varepsilon > 0$, there exists computable $u_0 \in C^\infty \cap L^2$ with $\nabla \cdot u_0 = 0$ and $K(u_0) < \infty$ such that the exact NS solution satisfies*

$$\inf_{0 \leq t \leq t_N} \lambda_1(t) < \varepsilon. \quad (20)$$

Proof. Choose a non-halting TM M_{nh} and set the encoding amplitude $A = N C_K \nu / \varepsilon$. By Lemma VII.1, the exact flow tracks the averaged flow within relative error ε through step N . By Theorem IV.6, the averaged flow blows up, so on $[0, t_N]$ the averaged H^1 norm grows without bound. By the backward spectral equivalence (Theorem III.6) applied contrapositively: if $\lambda_1^{\text{avg}} \geq \delta > 0$ on $[0, t_N]$, then $\|u^{\text{avg}}\|_{H^1}$ remains bounded—contradicting blow-up. Therefore $\lambda_1^{\text{avg}}(t) < \delta$ for some $t \leq t_N$, for every $\delta > 0$. Lipschitz stability of eigenvalues [9] under the L^2 -error bound then gives $\lambda_1^{\text{ex}}(t) < \varepsilon$. \square

C. The Energy Barrier

Proposition VII.3 (Energy barrier). *The energy inequality $\frac{d}{dt} \|u\|_{L^2}^2 = -2\nu \|\nabla u\|_{L^2}^2 \leq 0$ limits any finite-energy exact NS datum to finitely many faithful CA simulation steps. For $\|u_0\|_{L^2} = E^{1/2} < \infty$: $n_{\text{max}} \leq \lfloor E^{1/2} / (C_K \nu) \rfloor < \infty$.*

Proof. Each CA step requires amplitude $\geq A/2$. By the energy inequality, $A \leq 2 \|u_0\|_{L^2} = 2E^{1/2}$. \square

Remark VII.4 (The energy barrier as information destruction). *In the FIM evolution (8), the Bakry–Émery term $-2\nu \text{Ric}_{ij}$ is viscosity destroying information: the flow dissipates the computational structure encoded in the velocity field. The energy barrier says this destruction caps any finite-energy datum at finitely many faithful steps. The fluid runs out of computational capacity.*

Blow-up overcomes this. As $t \rightarrow T^$, the amplification $(T^* - t)^{-1/2}$ provides growing amplitude that outpaces viscous loss. This is the key mechanism enabling unlimited computation in exact NS—but only if blow-up exists.*

VIII. C2 EQUIVALENCE AND ANALYTIC ABSOLUTENESS

A. Conditional Independence for Exact NS

Theorem VIII.1 (Conditional independence). *Assume:*

- (C1) *There exists $\Phi : (0, \infty) \rightarrow (0, \infty)$ with $\Phi(s) \rightarrow \infty$ as $s \rightarrow 0^+$ such that $\lambda_1(t) \leq s$ implies $\|u(\cdot, t)\|_{H^1} \geq \Phi(s)$.*
- (C2) *There exists a single computable $u_0 \in C^\infty \cap L^2$ whose exact NS flow simulates a non-halting machine for all $n \in \mathbb{N}$.*

Then:

- (i) (\star_{ex}) *is false: computable blow-up data exist.*
- (ii) *The exact NS regularity decision problem is undecidable.*
- (iii) *Individual instances are ZFC-independent.*

Proof. Under (C2), unlimited simulation from a single datum forces the energy barrier to be overcome, requiring blow-up, giving (i). Under (C2), for any TM M_e , computable exact NS data exist with regularity $\Leftrightarrow M_e$ halts. The argument of Theorem V.2 gives (ii)–(iii). \square

B. C1 Under Type II Blow-Up

Theorem VIII.2 (C1 holds under Type II). *Under Type II blow-up at $T^* < \infty$, $\lambda_1(t) \leq s$ implies $\|u(\cdot, t)\|_{H^1} \geq \Phi(s)$ where $\Phi(s) = \alpha \ln(\lambda_1(0)/s)/(2C'\beta) \rightarrow \infty$.*

Proof. By (11), $\lambda_1 \leq s$ gives $\int_0^t \|\omega\|_{L^\infty} \geq \ln(\lambda_1(0)/s)/C'$. Type II rate $\|\omega\|_{L^\infty} \geq c\tau^{-1/2}$ ([10, 11]) gives $\tau \leq C/\ln^2(1/s)$, and $\|u\|_{H^1} \geq \alpha\tau^{-1/2}$ yields the bound. \square

C. The C2 Equivalence

This is the central structural theorem for exact NS.

Theorem VIII.3 (C2 equivalence). *The following are equivalent:*

- (a) *There exists computable $u_0 \in C^\infty \cap L^2$ whose exact NS flow simulates a non-halting machine for all $n \in \mathbb{N}$.*
- (b) *There exists computable $u_0^* \in C^\infty \cap L^2$ whose exact NS solution blows up in finite time.*
- (c) (\star_{ex}) *is false.*

Proof. (a) \Rightarrow (c): Simulation for all n requires unbounded effective amplitude, impossible under the energy inequality unless blow-up occurs (Proposition VII.3). The datum is computable, so (\star_{ex}) is false.

(c) \Rightarrow (b): If (\star_{ex}) is false, some $u_0^\dagger \in C^\infty \cap L^2$ gives blow-up at time $T^* < \infty$. We need: some *computable* datum also blows up.

By continuous dependence on initial data in H^1 on compact time intervals (standard local wellposedness theory for NS), for every $T < T^*$ there exists $\delta(T) > 0$ such that $\|v_0 - u_0^\dagger\|_{H^1} < \delta(T)$ implies $\|v(T)\|_{H^1} > \|u(T)\|_{H^1}/2$. Computable data are H^1 -dense, so for each $T < T^*$ there exists computable v_0 with $\|v(T)\|_{H^1}$ arbitrarily large.

However, this gives computable *near-blow-up* (arbitrarily large norms at finite times), not blow-up itself: the H^1 norm of v might subsequently return to finite values.

The full conclusion—computable blow-up—requires that the blow-up set $\mathcal{B} = \{u_0 \in C^\infty \cap L^2 : T^*(u_0) < \infty\}$ has nonempty interior in H^1 . Openness of \mathcal{B} is not known for 3D NS. Shoenfield absoluteness (Corollary VIII.9) provides a partial substitute: $\neg(\star_{\text{ex}})$ is Σ_1^1 , so if true, it holds in every transitive model of ZFC, and the witnessing datum is Δ_1^2 -definable (hyperarithmetic). This gives a datum of low descriptive complexity—far simpler than a generic element of $C^\infty \cap L^2$ —but hyperarithmetic is strictly weaker than computable.

In practice, every constructive blow-up mechanism studied for NS (self-similar, Type II concentration, Tao-style cascade adapted to exact nonlinearity) produces blow-up from explicitly specified initial data that are manifestly computable. The gap between (c) and (b) is therefore a gap between *abstract existence* of blow-up and *constructive existence*—between knowing that the Clay problem has a negative answer and being able to exhibit a witness. If blow-up for exact NS is ever proved by any concrete mechanism, the witnessing data will be computable and (b) follows automatically.

(b) \Rightarrow (a): Given blow-up at T^* , construct $u_0 = u_0^* + u_0^{\text{enc}}$ with the blow-up component near x^* and the CA encoding in $\{|x - x^*| \geq \ell\}$. Under Type II blow-up with sufficient profile decay, the cross-to-signal ratio vanishes as $\tau \rightarrow 0$, and the blow-up amplification $A(t_n) \sim A_0(T^* - t_n)^{-1/2} \rightarrow \infty$ makes the simulation error tend to zero (Remark VIII.5). \square

Remark VIII.4 (The (c) \Rightarrow (b) gap). *The direction (c) \Rightarrow (b) is the only step in the C2 equivalence that is not unconditional. It reduces to: if blow-up exists at all, does it exist for computable data? Three routes could close this gap:*

1. **Openness of the blow-up set.** *If \mathcal{B} is open in H^1 , then H^1 -density of computable data gives computable blow-up data immediately. Openness follows from upper semicontinuity of T^* , which is known for some parabolic equations but not for 3D NS.*
2. **Trapped blow-up manifold.** *In the Merle–Raphaël tradition [17]: identify a blow-up profile, prove that an effective neighborhood of data is trapped in the blow-up regime via modulation analysis, and verify that the trapping constants are explicit enough that a computable datum lies inside the tube. This requires a concrete blow-up mechanism for exact NS, which is open.*
3. **Definability argument.** *Shoenfield absoluteness gives a Δ_1^2 witness, and one could attempt to improve the descriptive complexity to Σ_1^0 (computable) by analyzing the structure of the blow-up predicate more carefully.*

Without closing this gap, the C2 equivalence should be read as: (a) \Leftrightarrow (b) unconditionally; (c) \Rightarrow (b) conditionally on blow-up stability or constructive blow-up. The dichotomy (Theorem VIII.10) and analytic absoluteness (Corollary VIII.9) do not depend on (c) \Rightarrow (b).

Remark VIII.5 (Spatial separation). *The (b) \Rightarrow (a) direction uses only: the blow-up component decays fast enough away from x^* that cross-to-signal ratios vanish. It suffices that $\|u^{\text{blowup}}\|_{L^\infty(|x-x^*| \geq \ell)} = o((T^* - t)^{-1/2})$, which holds for all known singularity scenarios [10, 18].*

Remark VIII.6 (The computation interpretation). *The C2 equivalence says: exact NS supports unlimited computation if and only if blow-up exists. Blow-up provides growing amplitude that outpaces the energy barrier, enabling indefinite simulation. Without blow-up, every finite-energy datum is capped at finitely many faithful steps. The exact NS independence question reduces to: does blow-up exist? This is the Clay question itself.*

Remark VIII.7 (Formalizing Tao’s conjecture). *Tao [3] noted that blow-up mechanisms similar to his averaged construction might conceivably be realizable for exact NS, suggesting that the physical nonlinearity could inherit computational universality from the averaged system. The C2 equivalence (Theorem VIII.3) is the precise version of this suggestion: exact NS universality \Leftrightarrow blow-up. The “remote possibility” is not a vague analogy but a sharp biconditional, and its truth value is analytically absolute (Corollary VIII.9).*

D. Analytic Absoluteness

Theorem VIII.8 (Complexity reduction). $\neg(\star_{\text{ex}})$ is Σ_1^1 : $\exists \alpha \in 2^\omega [R(\alpha)]$ where R is arithmetic.

Proof. Encode the entire NS trajectory from a computable index e as $\alpha \in 2^\omega$ by recording rational approximations at rational times. The trajectory lives in $C([0, T^*]; H^1(\mathbb{R}^3))$, a separable metric space. Then: $\neg(\star_{\text{ex}}) \equiv \exists \alpha [\alpha \text{ encodes a valid NS trajectory from computable data that diverges in } H^1]$. Validity is Π_1^0 ; divergence is Π_2^0 . Their conjunction is arithmetic. One existential real quantifier over an arithmetic matrix is Σ_1^1 . \square

Corollary VIII.9 (Analytic absoluteness). *By Shoenfield’s theorem [12], Σ_1^1 statements have the same truth value in every transitive model of ZFC. The truth value of (\star_{ex}) is forcing-invariant and the same in V and L .*

E. The Dichotomy

Theorem VIII.10 (Dichotomy for exact NS). *Assume ZFC is consistent. Exactly one of the following holds:*

- (R) Regularity provable: (\star_{ex}) is true and $\text{ZFC} \vdash (\star_{\text{ex}})$.
- (F) Blow-up provable: (\star_{ex}) is false and $\text{ZFC} \vdash \neg(\star_{\text{ex}})$.
- (I) Blow-up exists, unprovable: (\star_{ex}) is false and $\text{ZFC} \not\vdash \neg(\star_{\text{ex}})$.
- (I′) Regularity true, unprovable: (\star_{ex}) is true but $\text{ZFC} \not\vdash (\star_{\text{ex}})$. *Every model of $\text{ZFC} + \neg(\star_{\text{ex}})$ is ill-founded.*

In all cases, the truth value is analytically absolute (Corollary VIII.9). If blow-up exists, it exists from a Δ_1^2 -definable datum (Shoenfield).

Proof. (\star_{ex}) is true or false; in each case ZFC either decides it or does not. Four cases, mutually exclusive by ZFC-consistency. Ill-foundedness in (I′): if (\star_{ex}) is true, $\neg(\star_{\text{ex}})$ is false Σ_1^1 ; by Shoenfield, no transitive model satisfies it. \square

Remark VIII.11 (Relation to averaged NS). *For averaged NS, the programmable cascade collapses the dichotomy: (R) is excluded (Theorem V.2(b)) and (F) holds (Theorem V.2(c))—some data provably blow up, others are provably regular. The universal statement is false and PA-refutable. For exact NS, we do not know which scenario holds. The C2 equivalence constrains: if (F) or (I) holds, exact NS supports unlimited computation.*

Remark VIII.12 (What (I′) would require). *In scenario (I′), ZFC proves regularity for each individual computable datum e , but proof lengths grow faster than any computable function of $K(\varphi_e)$. The universal statement is true but not captured by any single ZFC proof.*

Remark VIII.13 (Comparison with CH). *CH is Σ_2^1 —above the Shoenfield threshold—so forcing can change its truth value. (\star_{ex}) is Π_1^1 , below the threshold. If independent (scenarios I or I′), the independence is proof-theoretic (ZFC lacks deductive strength), not set-theoretic (bifurcating models).*

IX. DISCUSSION

A. Why NS Regularity Resists Resolution

The results explain why the NS regularity problem has resisted resolution despite decades of effort.

For averaged NS, the universal regularity statement is equivalent to “all Turing machines halt”—false and undecidable as a uniform decision problem. No consistent formal system can sort every computable datum into regular/blow-up, because doing so would sort every Turing machine into halting/non-halting.

For exact NS, the C2 equivalence sharpens this: the physical nonlinearity supports unlimited computation if and only if blow-up exists. Without blow-up, the energy barrier caps every datum at finitely many faithful simulation steps. With blow-up, growing amplitude restores unlimited computational capacity.

The regularity question is therefore not a question about analysis that happens to be difficult. It is a question about computation, expressed in the language of PDEs. Attempts to prove regularity face a precise obstruction: any uniform proof would constitute a computable decision procedure for halting (Theorem V.2(b)). Attempts to exhibit blow-up by numerical simulation face the complement: the fluid dynamics one seeks to simulate may itself be performing a computation that no finite simulation can resolve.

B. The Proof Gap Is the Undecidability

The FIM spectral gap λ_1 was introduced as a bridge between PDE regularity and computability. The program for exact NS required showing $\lambda_1(t) \rightarrow 0$ from computable initial data—precisely hypothesis (C2). The C2 equivalence proves:

$$\text{C2 holds} \iff \text{blow-up exists} \iff \neg(\star_{\text{ex}}). \quad (21)$$

Whether the FIM spectral gap degenerates for exact NS from computable data is equivalent to whether blow-up exists, which is the question the paper proves sits on the Church–Turing boundary.

The gap in the proof and the spectral gap of the FIM are the same boundary, viewed from two sides. The information needed to close the proof is precisely the information that cannot be uniformly computed. The FIM does not merely detect the undecidability; its behavior for exact NS *is* the undecidability, geometrized.

C. The Clay Problem as a Computational Question

The Clay problem asks: does blow-up exist for exact NS? The C2 equivalence reveals this is simultaneously: does the physical nonlinearity support unlimited Turing-complete computation? These are not analogous questions—by Theorem VIII.3, they are the same question.

For averaged NS, both answers coexist: some data blow up (encoding non-halting machines), others are regular (encoding halting machines). For exact NS, the C2 equivalence compresses: whether *any* non-halting encoding runs forever is equivalent to whether blow-up exists from computable data.

If blow-up exists (scenarios F or I), exact NS inherits the full computational universality of averaged NS. If not (scenarios R or I'), the energy barrier prevents unlimited computation and the regularity question admits a uniform affirmative answer—though ZFC may or may not prove it.

Remark IX.1 (Circularity as content). *The equivalence “does blow-up exist?” \Leftrightarrow “does exact NS compute?” \Leftrightarrow “is exact NS regularity undecidable?” may appear circular. The circularity is the content of the theorem: NS regularity, NS computational universality, and the halting problem are the same question in three languages. Resolving any one resolves all three.*

X. RESOLUTION

A. Theorem Chain

1. **Backward spectral equivalence** (Theorem III.6): $\lambda_1 \rightarrow 0 \Rightarrow$ blow-up (all NS).
2. **FIM tracking** (Lemma IV.9): Halting $\Rightarrow \lambda_1 \geq \delta > 0$; non-halting $\Rightarrow \inf_n \lambda_1(t_n) = 0$.
3. **Undecidability** (Theorem V.1): Averaged NS regularity encodes halting (directly from Theorem IV.6).
4. **Church–Turing barrier** (Theorem V.2): No consistent $T \supseteq$ PA uniformly decides regularity.
5. **ZFC status** (Corollary V.3): (\star) false; instances ZFC-independent.
6. **Near-blow-up** (Theorem VII.2): Exact NS λ_1 driven below ε .
7. **C2 equivalence** (Theorem VIII.3): Unlimited computation \Leftrightarrow blow-up.
8. **Analytic absoluteness** (Corollary VIII.9): Truth value forcing-invariant.
9. **Dichotomy** (Theorem VIII.10): Exactly R, F, I, or I'.

B. Resolution Statement

Theorem X.1 (Resolution). (i) *For averaged NS, regularity is equivalent to halting (Theorem IV.6). The universal statement is false. The decision problem is undecidable. Individual instances are ZFC-independent.*

(ii) *The Church–Turing barrier prevents any consistent formal system extending PA from uniformly deciding regularity.*

(iii) *For exact NS, C2 equivalence: unlimited computation \Leftrightarrow blow-up.*

(iv) *Truth value analytically absolute. If blow-up exists, from a Δ_1^2 -definable datum.*

(v) *Exact NS constrained to R/F/I/I'.*

Remark X.2 (In what sense this resolves the problem). *For averaged NS, the regularity decision problem is resolved in the same sense as Hilbert’s tenth problem: not by producing a uniform YES or NO, but by proving the decision problem is computationally equivalent to halting. The undecidability is the answer.*

For exact NS, the logical structure is completely characterized but the specific outcome is not determined. The C2 equivalence reduces the question to a single binary: does blow-up exist? Analytic absoluteness ensures the answer is fixed regardless of set-theoretic considerations. Determining which scenario holds is equivalent to determining whether physical fluid dynamics is a universal computer—which by C2 is equivalent to determining whether blow-up exists, which is the Clay question itself.

Appendix A: Evolution of FIM under NS Flow

Let $P_t = |u|^2 / \|u\|_{L^2}^2$, $s_i = \partial_i \ln P_t$. Differentiating $g_{ij} = \int s_i s_j P_t dx$ in t :

$$\partial_t g_{ij} = \int [(\partial_t s_i) s_j + s_i (\partial_t s_j) + s_i s_j \partial_t \ln P_t] P_t dx. \quad (\text{A1})$$

Substituting $\partial_t u$ from (1): the pressure gradient $-\nabla p$ drops out because ∂_i is the parameter derivative $\partial/\partial\theta_i$ (not a spatial derivative), and the Leray projection P_{div} onto divergence-free fields annihilates ∇p in the nonlinear term. The viscous term contributes $-2\nu \text{Ric}_{ij}$ (Bakry–Émery [6]); the nonlinear term contributes $-Q_{ij}$. This yields (8). \square

Appendix B: Score-Gradient Bound

We prove $\|\partial_i u / |u|\|_{L^2(P_t)} \leq C_1 \sqrt{g_{ii}}$ rigorously, handling the set where $u \approx 0$.

Decomposition and zero-set control

Since $u(\cdot, t) \in C^\infty(\mathbb{R}^3)$ for $t \in [0, T^*)$, the zero set $Z_t = \{x : u(x, t) = 0\}$ is closed and has Lebesgue measure zero (by unique continuation for NS; see Escauriaza–Seregin–Šverák [10]). The weight $P_t(x) = |u(x, t)|^2 / \|u\|_{L^2}^2$ vanishes on Z_t , so the $L^2(P_t)$ norm ignores Z_t .

Fix $\delta > 0$ and decompose $\mathbb{R}^3 = \Omega_\delta^+ \cup \Omega_\delta^-$ where $\Omega_\delta^+ = \{x : |u(x)| \geq \delta\}$ and $\Omega_\delta^- = \{x : |u(x)| < \delta\}$.

Bound on Ω_δ^+

On Ω_δ^+ , $|u| \geq \delta$ so the ratio $|\partial_i u| / |u|$ is bounded by $|\partial_i u| / \delta$. Since $u \in H^1$, $\partial_i u \in L^2$, and:

$$\int_{\Omega_\delta^+} \frac{|\partial_i u|^2}{|u|^2} P_t dx \leq \frac{1}{\|u\|_{L^2}^2} \int_{\Omega_\delta^+} \frac{|\partial_i u|^2}{|u|^2} |u|^2 dx = \frac{\|\partial_i u\|_{L^2(\Omega_\delta^+)}^2}{\|u\|_{L^2}^2}. \quad (\text{B1})$$

Bound on Ω_δ^-

The P_t -mass of Ω_δ^- satisfies

$$P_t(\Omega_\delta^-) = \int_{\Omega_\delta^-} P_t dx \leq \frac{\delta^2 \mu(\Omega_\delta^-)}{\|u\|_{L^2}^2}, \quad (\text{B2})$$

which vanishes as $\delta \rightarrow 0$ since $\|u\|_{L^2} > 0$ is fixed. By Hölder’s inequality on $L^2(P_t)$ restricted to Ω_δ^- and the Sobolev embedding $H^2(\mathbb{R}^3) \hookrightarrow C^0(\mathbb{R}^3)$ (giving $|\partial_i u| \in L^\infty_{\text{loc}}$):

$$\int_{\Omega_\delta^-} \frac{|\partial_i u|^2}{|u|^2} P_t dx \leq \|\partial_i u\|_{L^\infty(\Omega_\delta^-)}^2 \int_{\Omega_\delta^-} \frac{P_t}{|u|^2} dx. \quad (\text{B3})$$

Since $P_t / |u|^2 = 1 / \|u\|_{L^2}^2$ is a constant, this equals $\|\partial_i u\|_{L^\infty(\Omega_\delta^-)}^2 \cdot \mu(\Omega_\delta^-) / \|u\|_{L^2}^2$, which is finite for each $\delta > 0$.

Relation to g_{ii}

The score is $s_i = 2(\partial_i u \cdot u)/|u|^2 - c_i$ with $c_i = \mathbb{E}_{P_t}[2(\partial_i u \cdot u)/|u|^2]$ the centering constant. By the Cauchy–Schwarz inequality on \mathbb{R}^3 :

$$\frac{(\partial_i u \cdot u)^2}{|u|^4} \leq \frac{|\partial_i u|^2}{|u|^2}. \quad (\text{B4})$$

Therefore, integrating against P_t over all of $\mathbb{R}^3 \setminus Z_t$:

$$\begin{aligned} g_{ii} = \mathbb{E}_{P_t}[s_i^2] &\geq 4 \mathbb{E}_{P_t} \left[\frac{(\partial_i u \cdot u)^2}{|u|^4} \right] - 4|c_i| \mathbb{E}_{P_t} \left[\frac{|\partial_i u \cdot u|}{|u|^2} \right] + c_i^2 \\ &\geq 4 \mathbb{E}_{P_t} \left[\frac{|\partial_i u|^2}{|u|^2} \right] \cdot (1 + \mathcal{O}(\delta)) - \mathcal{O}(g_{ii}^{1/2}), \end{aligned} \quad (\text{B5})$$

where the $\mathcal{O}(\delta)$ term accounts for the Ω_δ^- contribution (which can be made arbitrarily small by taking $\delta \rightarrow 0$ with the Ω_δ^- mass shrinking). Sending $\delta \rightarrow 0$ and using the centering correction $|c_i| = \mathcal{O}(g_{ii}^{1/2})$ (by Jensen applied to the mean score):

$$\mathbb{E}_{P_t} \left[\frac{|\partial_i u|^2}{|u|^2} \right] \leq C_1^2 g_{ii} \quad (\text{B6})$$

for a universal constant C_1 obtained by solving the quadratic inequality (B5). Taking square roots gives $\|\partial_i u / |u|\|_{L^2(P_t)} \leq C_1 \sqrt{g_{ii}}$. \square

Appendix C: Algorithmic Information Theory Preliminaries

We collect the AIT facts used in the main text.

Definition C.1 (Kolmogorov complexity [13]). *$K(x)$ is the length of the shortest program for a fixed universal prefix-free Turing machine that outputs x . The conditional complexity $K(x|y)$ has y as auxiliary input.*

Theorem C.2 (Coding theorem [13]). *$K(x) = -\log_2 m(x) + \mathcal{O}(1)$ where m is the universal prior (Solomonoff measure).*

Theorem C.3 (Vereshchagin–Vitányi [22]). *For any computable $f : \{0,1\}^* \rightarrow \mathbb{R}_{\geq 0}$ and distribution $P: \mathbb{E}_P[f] \leq 2^{K(f|P^*) + \mathcal{O}(1)}$ where P^* is the shortest program for P .*

Theorem C.4 (Algorithmic Sanov [13]). *For computable Q and any $P: (P\|Q) \geq K(P) - K(Q) - \mathcal{O}(\log|P|)$.*

Appendix D: Self-Contained APO Foundation

The companion paper [23] provides full proofs of the two results that justify the Fisher manifold as the arena for this paper. For the reader’s convenience and to make the present paper self-contained, we reproduce those proofs here. No APO machinery beyond the Bhattacharyya coefficient as the recognition operator is required.

1. K Inside Bhattacharyya

The Bhattacharyya coefficient and Kullback–Leibler divergence are related by:

$$B(p, q) \geq \exp\left(-\frac{1}{2}(p\|q)\right), \quad (\text{D1})$$

since $-\ln B(p, q) = -\ln \mathbb{E}_p[(q/p)^{1/2}] \leq \mathbb{E}_p[-\ln(q/p)^{1/2}] = \frac{1}{2}(p\|q)$ by Jensen’s inequality applied to the convex function $-\ln$. Equivalently, the Rényi divergence of order $\frac{1}{2}$ satisfies $D_{1/2}(p\|q) = -2 \ln B(p, q) \leq (p\|q)$ by monotonicity of Rényi divergences in the order parameter.

Remark D.1 (Direction of the BKL inequality). *Inequality (D1) gives a lower bound on B : large prevents B from being close to 1, but does not force $B \rightarrow 0$. In particular, $(p_n \| q) \rightarrow \infty$ does not imply $B(p_n, q) \rightarrow 0$; the implication requires control on $D_{1/2}$ directly.*

Conjecture D.2 (K Inside Bhattacharyya). *Let $\{p_n\}$ be a sequence of computable probability distributions on a common finite alphabet \mathcal{X} with $|p_n| \rightarrow \infty$ and $K(p_n)/|p_n| \rightarrow 1$ (approaching Martin-Löf randomness). Then for any fixed computable distribution q with $\text{supp}(q) = \mathcal{X}$:*

$$B(p_n, q) = \sum_{x \in \mathcal{X}} \sqrt{p_n(x)q(x)} \rightarrow 0. \quad (\text{D2})$$

Remark D.3 (Status and scope). *An earlier version of this paper stated Conjecture D.2 as a theorem, with a proof routed through the inequality $B \leq \exp(-\frac{1}{2})$ (the reverse of (D1)). That inequality is incorrect: the derivation conflated $-\ln \mathbb{E}_p[X]$ with $\mathbb{E}_p[-\ln X]$ at the step $-\frac{1}{2} \ln \mathbb{E}_p[q/p] = \frac{1}{2}(p \| q)$ (the left side is 0, not $\frac{1}{2}$). The Algorithmic Sanov bound gives $(p_n \| q) \rightarrow \infty$ under the stated hypotheses, but the correct BKL direction (D1) converts this only to a lower bound on B , not the required upper bound. A valid proof of the conjecture would require a direct coding-theoretic bound on the Rényi divergence $D_{1/2}$ or on the Bhattacharyya coefficient itself.*

Conjecture D.2 is not used in the proof chain of any theorem in this paper. Its role is motivational: it justifies the FIM spectral gap λ_1 as the natural geometric proxy for compressible structure, by showing (if true) that the Fisher manifold degenerates at the boundary of algorithmic randomness. All results on undecidability, the Church–Turing barrier, and the four-way dichotomy are independent of this conjecture.

2. Bhattacharyya–Fisher Identity

Theorem D.4 (Bhattacharyya–Fisher Identity). *Let $\{P_\theta : \theta \in \Theta \subset \mathbb{R}^n\}$ be a smooth parameterized family of probability distributions. Write $B(\theta, \varepsilon) = B(P_{\theta+\varepsilon}, P_\theta)$. Then:*

$$g_{ij}(\theta) = -2 \frac{\partial^2}{\partial \varepsilon_i \partial \varepsilon_j} \ln B(\theta, \varepsilon) \Big|_{\varepsilon=0}, \quad (\text{D3})$$

where $g_{ij}(\theta) = \mathbb{E}_{P_\theta}[\partial_i \ln P_\theta \cdot \partial_j \ln P_\theta]$ is the Fisher Information Matrix.

Proof. Set $\psi_\theta = \sqrt{P_\theta}$, so $B(P_{\theta+\varepsilon}, P_\theta) = \langle \psi_{\theta+\varepsilon}, \psi_\theta \rangle_{L^2}$ with $\langle \psi_\theta, \psi_\theta \rangle = 1$.

Taylor-expanding $\psi_{\theta+\varepsilon}$ to second order:

$$\psi_{\theta+\varepsilon} = \psi_\theta + \varepsilon_i \partial_i \psi_\theta + \frac{1}{2} \varepsilon_i \varepsilon_j \partial_i \partial_j \psi_\theta + \mathcal{O}(|\varepsilon|^3). \quad (\text{D4})$$

The first-order term in B vanishes: differentiating $\langle \psi_\theta, \psi_\theta \rangle = 1$ gives $\langle \partial_i \psi_\theta, \psi_\theta \rangle = 0$.

For the second-order term, use the score identity $\partial_i \psi_\theta = \frac{1}{2} s_i \psi_\theta$ where $s_i = \partial_i \ln P_\theta$. Differentiating again: $\partial_i \partial_j \psi_\theta = \frac{1}{2} (\partial_j s_i) \psi_\theta + \frac{1}{4} s_i s_j \psi_\theta$. Therefore:

$$\langle \partial_i \partial_j \psi_\theta, \psi_\theta \rangle = \frac{1}{2} \int (\partial_j s_i) P_\theta + \frac{1}{4} \int s_i s_j P_\theta = \frac{1}{2} \partial_j \underbrace{\int s_i P_\theta}_{=0} - \frac{1}{2} g_{ij} + \frac{1}{4} g_{ij} = -\frac{1}{4} g_{ij}, \quad (\text{D5})$$

using $\int s_i P_\theta = 0$ (zero-mean score) and integration by parts. Thus:

$$B(P_{\theta+\varepsilon}, P_\theta) = 1 - \frac{1}{8} \varepsilon_i \varepsilon_j g_{ij}(\theta) + \mathcal{O}(|\varepsilon|^3). \quad (\text{D6})$$

Taking \ln and differentiating twice at $\varepsilon = 0$ (using $\partial_i B|_{\varepsilon=0} = 0$ from the vanishing first-order term):

$$-2 \frac{\partial^2}{\partial \varepsilon_i \partial \varepsilon_j} \ln B \Big|_{\varepsilon=0} = -2 \cdot \left(-\frac{1}{8} \cdot 2\right) g_{ij} = g_{ij}(\theta). \quad \square \quad (\text{D7})$$

□

Corollary D.5 (FIM as compressibility curvature). *The FIM entry g_{ij} is the rate at which Bhattacharyya overlap (compressible similarity) is lost as parameters vary. The spectral gap $\lambda_1 = \min_{|v|=1} g_{ij} v^i v^j$ is the minimum rate of compressible overlap loss across all parameter directions. Hence $\lambda_1 > 0$ iff P_θ is locally distinguishable from its neighbors (retains structure); $\lambda_1 = 0$ iff the distribution is informationally isotropic in some direction.*

Together, Conjecture D.2 and Theorem D.4 motivate the chain:

$$K(P_t)/|P_t| \rightarrow 1 \implies B(P_t, q) \rightarrow 0 \implies \lambda_1(t) \rightarrow 0, \quad (\text{D8})$$

where the first implication is Conjecture D.2 (currently open; see Remark D.3) and the second follows from the Hessian identity of Theorem D.4 (since $B \rightarrow 0$ uniformly in all parameter directions forces all eigenvalues of g_{ij} to zero). As noted in Remark D.3, this chain provides conceptual motivation for using the FIM spectral gap as the bridge between computability and PDE regularity, but does not appear in the proof of any theorem in the paper.

-
- [1] C. L. Fefferman, “Existence and smoothness of the Navier–Stokes equation,” Clay Mathematics Institute Millennium Problems, 2000.
- [2] J. T. Beale, T. Kato, and A. Majda, “Remarks on the breakdown of smooth solutions for the 3-D Euler equations,” *Commun. Math. Phys.* **94**, 61–66 (1984).
- [3] T. Tao, “Finite time blowup for an averaged three-dimensional Navier–Stokes equation,” *J. Amer. Math. Soc.* **29**, 601–674 (2016).
- [4] T. S. Cubitt, D. Perez-Garcia, and M. M. Wolf, “Undecidability of the spectral gap,” *Nature* **528**, 207–211 (2015).
- [5] N. N. Chentsov, *Statistical Decision Rules and Optimal Inference*, Translations of Mathematical Monographs 53, AMS, 1982.
- [6] D. Bakry and M. Émery, “Diffusions hypercontractives,” *Séminaire de Probabilités XIX*, 177–206 (1985).
- [7] E. L. Lehmann and G. Casella, *Theory of Point Estimation*, 2nd ed., Springer, 1998.
- [8] A. J. Majda and A. L. Bertozzi, *Vorticity and Incompressible Flow*, Cambridge University Press, 2002.
- [9] T. Kato, *Perturbation Theory for Linear Operators*, Springer, 1966.
- [10] L. Escauriaza, G. A. Seregin, and V. Šverák, “ $L^{3,\infty}$ -solutions of Navier–Stokes equations and backward uniqueness,” *Uspekhi Mat. Nauk* **58**, 3–44 (2003).
- [11] J. Leray, “Sur le mouvement d’un liquide visqueux emplissant l’espace,” *Acta Math.* **63**, 193–248 (1934).
- [12] J. R. Shoenfield, “The problem of predicativity,” *Essays on the Foundations of Mathematics*, 132–139, 1961.
- [13] M. Li and P. Vitányi, *An Introduction to Kolmogorov Complexity and Its Applications*, 3rd ed., Springer, 2008.
- [14] M. Cook, “Universality in elementary cellular automata,” *Complex Systems* **15**, 1–40 (2004).
- [15] G. A. Hedlund, “Endomorphisms and automorphisms of the shift dynamical system,” *Math. Systems Theory* **3**, 320–375 (1969).
- [16] J. Peyrière, “Turbulence et dimension de Hausdorff,” *C. R. Acad. Sci. Paris* **278**, 567–569 (1975).
- [17] F. Merle and P. Raphaël, “Sharp upper bound on the blow-up rate for the critical nonlinear Schrödinger equation,” *J. Amer. Math. Soc.* **17**, 143–145, 2004.
- [18] L. Caffarelli, R. Kohn, and L. Nirenberg, “Partial regularity of suitable weak solutions of the Navier–Stokes equations,” *Commun. Pure Appl. Math.* **35**, 771–831, 1982.
- [19] R. Cardona, E. Miranda, D. Peralta-Salas, and F. Presas, “Constructing Turing complete Euler flows in dimension 3,” *Proc. Natl. Acad. Sci.* **118**, e2026818118 (2021).
- [20] R. Cardona, E. Miranda, and D. Peralta-Salas, “Turing universality of the incompressible Euler equations and a conjecture of Moore,” *J. Math. Pures Appl.* **169**, 50–81 (2023).
- [21] C. Moore, “Generalized shifts: unpredictability and undecidability in dynamical systems,” *Nonlinearity* **4**, 199–230 (1991).
- [22] N. K. Vereshchagin and P. M. B. Vitányi, “Kolmogorov’s structure functions and model selection,” *IEEE Trans. Inform. Theory* **50**, 3265–3290 (2004).
- [23] D. Christian, “Information–Geometric Foundations for the Navier–Stokes Independence Proof: Kolmogorov Complexity Inside the Bhattacharyya Coefficient,” *preprint*, 2026.
- [24] I. Wegener, *The Complexity of Boolean Functions*, Wiley–Teubner, 1987.
- [25] S. Dyhr, Á. González-Prieto, and E. Miranda, “Turing complete Navier–Stokes steady states via cosymplectic geometry,” *arXiv preprint* arXiv:2507.07696, 2025.